

Fusion for Free

Efficient Algebraic Effect Handlers

Nicolas Wu¹ and Tom Schrijvers²

¹ Department of Computer Science, University of Bristol

² Department of Computer Science, KU Leuven

Abstract. Algebraic effect handlers are a recently popular approach for modelling side-effects that separates the syntax and semantics of effectful operations. The shape of syntax is captured by functors, and free monads over these functors denote syntax trees. The semantics is captured by algebras, and effect handlers pass these over the syntax trees to interpret them into a semantic domain.

This approach is inherently modular: different functors can be composed to make trees with richer structure. Such trees are interpreted by applying several handlers in sequence, each removing the syntactic constructs it recognizes. Unfortunately, the construction and traversal of intermediate trees is painfully inefficient and has hindered the adoption of the handler approach.

This paper explains how a sequence of handlers can be fused into one, so that multiple tree traversals can be reduced to a single one and no intermediate trees need to be allocated. At the heart of this optimization is keeping the notion of a free monad abstract, thus enabling a change of representation that opens up the possibility of fusion. We demonstrate how the ensuing code can be inlined at compile time to produce efficient handlers.

1 Introduction

Free monads are currently receiving a lot of attention. They are at the heart of *algebraic effect handlers*, a new purely functional approach for modelling side effects introduced by Plotkin and Power [15]. Much of their appeal stems from the separation of the syntax and semantics of effectful operations. This is both conceptually simple and flexible, as multiple different semantics can be provided for the same syntax.

The syntax of the primitive side-effect operations is captured in a signature functor. The free monad over this functor assembles the syntax for the individual operations into an abstract syntax tree for an effectful program. The semantics of the individual operations is captured in an algebra, and an effect handler folds the algebra over the syntax tree of the program to interpret it into a semantic domain.

A particular strength of the approach is its inherent modularity. Different signature functors can be composed to make trees with richer structure. Such

trees are interpreted by applying several handlers in sequence, each removing the syntactic constructs it recognizes.

Unfortunately, the construction and traversal of intermediate trees is rather costly. This inefficiency is perceived as a serious weakness of effect handlers, especially when compared to the traditional approach of composing effects with monad transformers. While several authors address the cost of constructing syntax trees with free monads, efficiently applying multiple handlers in sequence has received very little attention. As far as we know, only Kammar et al. [9] provide an efficient implementation. Unfortunately, this implementation does not come with an explanation. Hence it is underappreciated and ill-understood.

In this paper we close the gap and explain how a sequence of algebraic effect handlers can be effectively fused into a single handler. Central to the paper are the many interpretations of the word *free*. Interpreting *free* monads as the initial objects of the more general term algebras and, in particular, term monads provide an essential change of perspective where *free* theorems enable fusion. The codensity monad facilitates the way, turning any term algebra into a monadic one for *free*, and with an appropriate code setup in Haskell the GHC compiler takes care of fusion at virtually no cost³ to the programmer. The result is an effective implementation that compares well to monad transformers.

2 Algebraic Effect Handlers

The idea of the algebraic effect handlers approach is to consider the free monad over a particular functor as an abstract syntax tree (AST) for an effectful computation. The functor is used to generate the nodes of a free structure whose leaves correspond to variables. This can be defined as an inductive datatype $Free\ f$ for a given functor f .

```
data Free f a where
  Var :: a → Free f a
  Con :: f (Free f a) → Free f a
```

The nodes are *constructed* by Con , and the *variables* are given by Var .

Since a value of type $Free\ f\ a$ is an inductive structure, we can define a *fold* for it by providing a function gen that deals with generation of values from $Var\ x$, and an algebra alg that is used to recursively collapse an operation $Con\ op$.

```
fold :: Functor f ⇒ (f b → b) → (a → b) → (Free f a → b)
fold alg gen (Var x) = gen x
fold alg gen (Con op) = alg (fmap (fold alg gen) op)
```

Algebraic effect handlers give a semantics to the syntax tree: one way of doing this is by using a *fold*.

³ Yes, almost for *free*!

The behaviour of folds when composed with other functions is described by fusion laws. The first law describes how certain functions that are precomposed with a fold can be incorporated into a new fold:

$$\text{fold } \text{alg } \text{gen} \cdot \text{fmap } h = \text{fold } \text{alg } (\text{gen} \cdot h) \quad (1)$$

The second law shows how certain functions that are postcomposed with a fold can be incorporated into a new fold:

$$k \cdot \text{fold } \text{alg } \text{gen} = \text{fold } \text{alg}' (k \cdot \text{gen}) \quad (2)$$

this is subject to the condition that $k \cdot \text{alg} = \text{alg}' \cdot \text{fmap } k$.

The monadic instance of the free monad is witnessed by the following:

```
instance Functor f  $\Rightarrow$  Monad (Free f) where
  return x = Var x
  m  $\gg=$  f = fold Con f m
```

Variables are the way of providing a return for the monad, and extending a syntax tree by means of a function f corresponds to applying that function to the variables found at the leaves of the tree.

2.1 Nondeterminism

A functor supplies the abstract syntax for the primitive effectful operations in the free monad. For instance, the *Nondet* functor provides the *Or* k k syntax for a binary nondeterministic choice primitive. The parameter to the constructor of type k marks the recursive site of syntax, which indicates where the continuation is after this syntactic fragment has been evaluated.

```
data Nondet k where
  Or :: k  $\rightarrow$  k  $\rightarrow$  Nondet k
instance Functor Nondet where
  fmap f (Or x y) = Or (f x) (f y)
```

This allows us to express the syntax tree of a computation that nondeterministically returns *True* or *False*.

```
coin :: Free Nondet Bool
coin = Con (Or (Var True) (Var False))
```

The syntax is complemented by semantics in the form of effect handlers—functions that replace the syntax by values from a semantic domain. Using a *fold* for the free monad is a natural way of expressing such functions. For instance, here is a handler that interprets *Nondet* in terms of lists of possible outcomes:

```
handleNondet :: Free Nondet a  $\rightarrow$  [a]
handleNondet = fold algNondet genNondet
```

where $alg_{Nondet_{\square}}$ is the *Nondet*-algebra that interprets terms constructed by *Or* operations and $gen_{Nondet_{\square}}$ interprets variables.

$$\begin{aligned} alg_{Nondet_{\square}} &:: Nondet [a] \rightarrow [a] \\ alg_{Nondet_{\square}} (Or \ l_1 \ l_2) &= l_1 \uplus l_2 \\ gen_{Nondet_{\square}} &:: a \rightarrow [a] \\ gen_{Nondet_{\square}} x &= [x] \end{aligned}$$

The variables of the syntax tree are turned into singleton lists by $gen_{Nondet_{\square}}$, and choices between alternatives are put together by $alg_{Nondet_{\square}}$, which appends lists.

As an example, we can interpret the *coin* program:

$$\begin{aligned} &> handle_{Nondet_{\square}} coin \\ &[True, False] \end{aligned}$$

This particular interpretation gives us a list of the possible outcomes.

Generalizing away from the details, handlers are usually presented in the form

$$hdl :: \forall a . Free \ F \ a \rightarrow H \ a$$

where F and H are arbitrary functors determined by the handler.

2.2 Handler Composition

There are many useful scenarios that involve the (function) composition of effect handlers. We now consider two classes of this kind of composition.

Effect Composition A first important class of scenarios is where multiple effects are combined in the same program. To this end we compose signatures and handlers.

The coproduct functor $f + g$ makes it easy to compose functors:

$$\begin{aligned} \mathbf{data} \ (+) \ f \ g \ a \ \mathbf{where} \\ \quad Inl &:: f \ a \rightarrow (f + g) \ a \\ \quad Inr &:: g \ a \rightarrow (f + g) \ a \\ \mathbf{instance} \ (Functor \ f, Functor \ g) \Rightarrow Functor \ (f + g) \ \mathbf{where} \\ \quad fmap \ f \ (Inl \ s) &= Inl \ (fmap \ f \ s) \\ \quad fmap \ f \ (Inr \ s) &= Inr \ (fmap \ f \ s) \end{aligned}$$

The free monad of a coproduct functor is a tree where each node can be built from syntax from either f or g .

Composing handlers is easy too: if the handlers are written in a compositional style, then function composition does the trick. A compositional handler for the functor F has a signature of the form:

$$hdl :: \forall g \ a . Free \ (F + g) \ a \rightarrow H_1 \ (Free \ g \ (G_1 \ a))$$

This processes only the F -nodes in the AST and leaves the g -nodes as they are. Hence the result of the compositional handler is a new (typically smaller) AST with only g -nodes. The variables of type $G_1 a$ in the resulting AST are derived from the variables of type a in original AST as well as from the processed operations. Moreover, the new AST may be embedded in a context H_1 .

For instance, the compositional nondeterminism handler is defined as follows, where $F = \text{Nondet}$, $G_1 = []$ and implicitly $H_1 = \text{Id}$.

$$\begin{aligned} \text{handle}_{\text{Nondet}} &:: \text{Functor } g \Rightarrow \text{Free } (\text{Nondet} + g) a \rightarrow \text{Free } g [a] \\ \text{handle}_{\text{Nondet}} &= \text{fold } (\text{alg}_{\text{Nondet}} \nabla \text{Con}) \text{ gen}_{\text{Nondet}} \end{aligned}$$

Here the variables are handled with the monadified version of $\text{gen}_{\text{Nondet}[]}$, given by $\text{gen}_{\text{Nondet}}$.

$$\begin{aligned} \text{gen}_{\text{Nondet}} &:: \text{Functor } g \Rightarrow a \rightarrow \text{Free } g [a] \\ \text{gen}_{\text{Nondet}} x &= \text{Var } [x] \end{aligned}$$

The g nodes are handled by a Con algebra, which essentially leaves them untouched. The Nondet nodes are handled by the $\text{alg}_{\text{Nondet}}$ algebra, which is a monadified version of $\text{alg}_{\text{Nondet}[]}$.

$$\begin{aligned} \text{alg}_{\text{Nondet}} &:: \text{Functor } g \Rightarrow \text{Nondet } (\text{Free } g [a]) \rightarrow \text{Free } g [a] \\ \text{alg}_{\text{Nondet}} (\text{Or } ml_1 \text{ } ml_2) &= \\ \text{do } \{ l_1 \leftarrow ml_1; l_2 \leftarrow ml_2; \text{Var } (l_1 \text{ } \text{ } l_2) \} \end{aligned}$$

The junction combinator (∇) composes the algebras for the two kinds of nodes.

$$\begin{aligned} (\nabla) &:: (f \text{ } b \rightarrow b) \rightarrow (g \text{ } b \rightarrow b) \rightarrow ((f + g) \text{ } b \rightarrow b) \\ (\nabla) \text{ alg}_f \text{ alg}_g (\text{Inl } s) &= \text{alg}_f s \\ (\nabla) \text{ alg}_f \text{ alg}_g (\text{Inr } s) &= \text{alg}_g s \end{aligned}$$

In the definition of handlerNondet , we use $\text{alg}_{\text{Nondet}} \nabla \text{Con}$. Since the functor in question is $\text{Nondet} + g$, the values constructed by Nondet are handled by $\text{alg}_{\text{Nondet}}$, and values constructed by g are left untouched: the fold unwraps one level of Con , but then replaces it with a Con again.

A second example of an effect signature is that of state, whose primitive operations *Get* and *Put* respectively query and modify the implicit state.

$$\begin{aligned} &\text{data State } s \text{ } k \text{ where} \\ &\quad \text{Put} :: s \rightarrow k \rightarrow \text{State } s \text{ } k \\ &\quad \text{Get} :: (s \rightarrow k) \rightarrow \text{State } s \text{ } k \\ &\text{instance Functor (State } s) \text{ where} \\ &\quad \text{fmap } f (\text{Put } s \text{ } k) = \text{Put } s (f \text{ } k) \\ &\quad \text{fmap } f (\text{Get } k) = \text{Get } (f \cdot k) \end{aligned}$$

The compositional handler for state is as follows, where $F = \text{State } s$, $H_1 = s \rightarrow -$ and implicitly $G_1 = \text{Id}$.

$$\begin{aligned} \text{handle}_{\text{State}} &:: \text{Functor } g \Rightarrow \text{Free } (\text{State } s + g) \ a \rightarrow (s \rightarrow \text{Free } g \ a) \\ \text{handle}_{\text{State}} &= \text{fold } (\text{alg}_{\text{State}} \nabla \text{con}_{\text{State}}) \ \text{gen}_{\text{State}} \end{aligned}$$

This time the variable and constructor cases are defined as:

$$\begin{aligned} \text{gen}_{\text{State}} &:: \text{Functor } g \Rightarrow a \rightarrow (s \rightarrow \text{Free } g \ a) \\ \text{gen}_{\text{State}} \ x \ s &= \text{Var } x \\ \text{alg}_{\text{State}} &:: \text{Functor } g \Rightarrow \text{State } s \ (s \rightarrow \text{Free } g \ a) \rightarrow (s \rightarrow \text{Free } g \ a) \\ \text{alg}_{\text{State}} \ (\text{Put } s' \ k) \ s &= k \ s' \\ \text{alg}_{\text{State}} \ (\text{Get } k) \ s &= k \ s \ s \end{aligned}$$

Using $\text{gen}_{\text{State}}$, a variable x is replaced by a version that ignores any incoming state parameter s . Any stateful constructs are handled by $\text{alg}_{\text{State}}$, where a continuation k proceeds by using the appropriate state: if the syntax is a $\text{Put } s' \ k$, then the new state is s' , otherwise the syntax is $\text{Get } k$, in which case the state is left unchanged and passed as s .

Finally, any syntax provided by g is adapted by $\text{con}_{\text{State}}$ to take the extra state parameter s into account:

$$\begin{aligned} \text{con}_{\text{State}} &:: \text{Functor } g \Rightarrow g \ (s \rightarrow \text{Free } g \ a) \rightarrow (s \rightarrow \text{Free } g \ a) \\ \text{con}_{\text{State}} \ op \ s &= \text{Con } (\text{fmap } (\lambda m \rightarrow m \ s) \ op) \end{aligned}$$

This feeds the state s to the continuations of the operation.

To demonstrate effect composition we can put *Nondet* and *State* together and handle them both. Before we do so, we also need a base case for the composition, which is the empty signature *Void*.

```
data Void k
instance Functor Void
```

The *Void* handler only provides a variable case since the signature has no constructors. In fact, a *Free Void* term can only be a *Var x*, so x is immediately output using the identity function.

$$\begin{aligned} \text{handle}_{\text{Void}} &:: \text{Free } \text{Void} \ a \rightarrow a \\ \text{handle}_{\text{Void}} &= \text{fold } \perp \ \text{id} \end{aligned}$$

Finally, we can put together a composite handler for programs that feature both nondeterminism and state. The signature of such programs is the composition of the three basic signatures:

```
type  $\Sigma = \text{Nondet} + (\text{State } \text{Int} + \text{Void})$ 
```

The handler is the composition of the three handlers, working from the left-most functor in the signature:

$$\begin{aligned} \text{handle}_{\Sigma} &:: \text{Free } \Sigma \ a \rightarrow \text{Int} \rightarrow [a] \\ \text{handle}_{\Sigma} \ \text{prog} &= \text{handle}_{\text{Void}} \cdot (\text{handle}_{\text{State}} \cdot \text{handle}_{\text{Nondet}}) \ \text{prog} \end{aligned}$$

Effect Delegation Another important class of applications are those where a handler expresses the complex semantics of particular operations in terms of more primitive effects.

For instance, the following logging handler for state records every update of the state by means of the *Writer* effect.

$$\begin{aligned} \text{handle}_{\text{LogState}} &:: \text{Free} (\text{State } s) a \rightarrow s \rightarrow \text{Free} (\text{Writer String} + \text{Void}) a \\ \text{handle}_{\text{LogState}} &= \text{fold } \text{alg}_{\text{LogState}} \text{ gen}_{\text{State}} \\ \text{alg}_{\text{LogState}} &:: \text{State } s (s \rightarrow \text{Free} (\text{Writer String} + \text{Void}) a) \\ &\quad \rightarrow s \rightarrow \text{Free} (\text{Writer String} + \text{Void}) a \\ \text{alg}_{\text{LogState}} (\text{Put } s' k) s &= \text{Con} (\text{Inl} (\text{Tell "put" } (k s'))) \\ \text{alg}_{\text{LogState}} (\text{Get } k) s &= k s s \end{aligned}$$

The syntax of the *Writer* effect is captured by the following functor, where w is a parameter that represents the type of values that are written to the log:

```
data Writer w k where
  Tell :: w → k → Writer w k
instance Functor (Writer w) where
  fmap f (Tell w k) = Tell w (f k)
```

A semantics can be given by the following handler, where w is constrained to be a member of the *Monoid* typeclass.

$$\begin{aligned} \text{handle}_{\text{Writer}} &:: (\text{Functor } g, \text{Monoid } w) \Rightarrow \text{Free} (\text{Writer } w + g) a \rightarrow \text{Free } g (w, a) \\ \text{handle}_{\text{Writer}} &= \text{fold} (\text{alg}_{\text{Writer}} \nabla \text{Con}) \text{ gen}_{\text{Writer}} \end{aligned}$$

The variables are evaluated by pairing with the unit of the monoid given by *empty* before being embedded into the monad m_2 .

$$\begin{aligned} \text{gen}_{\text{Writer}} &:: (\text{Monad } m_2, \text{Monoid } w) \Rightarrow a \rightarrow m_2 (w, a) \\ \text{gen}_{\text{Writer}} x &= \text{return} (\text{empty}, x) \end{aligned}$$

When a *Tell* $w_1 k$ operation is encountered, the continuation k is followed by a state where w_1 is appended using *mappend* to any generated logs.

$$\begin{aligned} \text{alg}_{\text{Writer}} &:: (\text{Monad } m_2, \text{Monoid } w) \Rightarrow \text{Writer } w (m_2 (w, a)) \rightarrow m_2 (w, a) \\ \text{alg}_{\text{Writer}} (\text{Tell } w_1 k) &= k \gg= \lambda(w_2, x) \rightarrow \text{return} (w_1 \text{ 'mappend' } w_2, x) \end{aligned}$$

To see this machinery in action, consider the following program that makes use of state:

```
program :: Int → Free (State Int) Int
program n
  | n ≤ 0    = Con (Get var)
  | otherwise = Con (Get (λs → Con (Put (s + n) (program (n - 1)))))
```

This is then simply evaluated by running handlers in sequence.

$example :: Int \rightarrow (String, Int)$
 $example\ n = (handle_{Void} \cdot handle_{Writer} \cdot handle_{LogState}\ (program\ n))\ 0$

To fully interpret a stateful program, we must first run $handle_{LogState}$, which interprets the *Tell* operations by generating a tree with *Writer String* syntax. This generated syntax is then handled with the $handle_{Writer}$ handler.

3 Fusion

The previous composition examples lead us to the main challenge of this paper: The composition of two handlers produces an intermediate abstract syntax tree. How can we fuse the two handlers into a single one that does not involve an intermediate tree?

More concretely, given two handlers of the form:

$handler_1 :: Free\ F_1\ a \rightarrow H_1\ (G_1\ a)$
 $handler_1 = fold\ alg_1\ gen_1$
 $handler_2 :: Free\ F_2\ a \rightarrow H_2\ a$
 $handler_2 = fold\ alg_2\ gen_2$

where F_1 and F_2 are signature functors and H_1 , G_1 and H_2 are arbitrary functors, our goal is to obtain a combined handler

$pipeline_{12} :: Free\ F_1\ a \rightarrow H_1\ (H_2\ (G_1\ a))$
 $pipeline_{12} = fold\ alg_{12}\ gen_{12}$

such that

$$fmap\ handler_2 \cdot handler_1 = pipeline_{12} \quad (3)$$

3.1 Towards Proper Builders

The fact that $handler_1$ *builds* an AST over functor F_2 and that $handler_2$ *folds* over this AST suggests a particular kind of fusion known as shortcut fusion or fold/build fusion [4].

One of the two key ingredients for this kind of fusion is already manifestly present: the *fold* in $handler_2$. Yet, the structure and type of $handler_1$ do not necessarily force it to be a proper builder: fold/build fusion requires that the builder creates the F_2 -AST from scratch by generating all the *Var* and *Con* constructors itself. Indeed, in theory, $handler_1$ could produce the F_2 -AST out of ready-made components supplied by (fields of) a colluding F_1 functor.

In order to force $handler_1$ to be a proper builder, we require it to be implemented against a builder interface rather than a concrete representation. This builder interface is captured in the typeclass *TermMonad* (explained below), and then, with the following constraint polymorphic signature, $handler_1$ is guaranteed to build properly:

$handler_1 :: (TermMonad\ m_2\ F_2) \Rightarrow Free\ F_1\ a \rightarrow H_1\ (m_2\ (G_1\ a))$

Term Algebras The concept of a *term algebra* provides an abstract interface for the primitive ways to build an AST: the two constructors *Con* and *Var* of the free monad. We borrow the nomenclature from the literature on universal algebras [2].

A *term algebra* is an f -algebra $con :: f (h a) \rightarrow h a$ with a carrier $h a$. The values in $h a$ are those generated by the set of variables a with a valuation function $var :: a \rightarrow h a$, as well as those that arise out of repeated applications of the algebra. This is modelled by the typeclass *TermAlgebra* $h f$ as follows:

```
class Functor f  $\Rightarrow$  TermAlgebra h f | h  $\rightarrow$  f where
  var ::  $\forall a . a \rightarrow h a$ 
  con ::  $\forall a . f (h a) \rightarrow h a$ 
```

The function *var* is used to embed a *variable* into the term, and the function *con* is used to *construct* a term from existing ones. This typeclass is well-defined only when $h a$ is indeed generated by *var* and *con*.

The most trivial instance of this typeclass is of course that of the free monad.

```
instance Functor f  $\Rightarrow$  TermAlgebra (Free f) f where
  var = Var
  con = Con
```

Term Monads There are two additional convenient ways to build an AST: the monadic primitives *return* and (\gg).

A monad m is a *term monad* for a functor f , if there is a term algebra for f whose carrier is $m a$. We can model this relationship as a typeclass with no members.

```
class (Monad m, TermAlgebra m f)  $\Rightarrow$  TermMonad m f | m  $\rightarrow$  f
instance (Monad m, TermAlgebra m f)  $\Rightarrow$  TermMonad m f
```

Again, the free monad is the obvious instance of this typeclass.

```
instance Functor f  $\Rightarrow$  TermMonad (Free f) f
```

Its monadic primitives are implemented in terms of *fold*, *con* and *var*. In the abstract builder interface *TermMonad* we only partially expose this fact, by means of the following two laws. Firstly, the *var* operation should coincide with the monad's *return*.

$$var = return \tag{4}$$

Secondly, the monad's bind (\gg) should distribute through *con*.

$$con\ op\ \gg\ k = con\ (fmap\ (\gg\ k)\ op) \tag{5}$$

This law states that a term constructed by an operation, where the term is followed by a continuation k , is equivalent to constructing a term from an operation where the operation is followed by k . In other words, the arguments of an operation correspond to its continuations.

Examples All the compositional handlers of Section 2.2 can be easily expressed in terms of the more abstract *TermMonad* interface. For example, the revised nondeterminism handler looks as follows.

$$\begin{aligned}
& \text{handle}'_{\text{Nondet}} :: \text{TermMonad } m \ g \Rightarrow \text{Free } (\text{Nondet} + g) \ a \rightarrow m \ [a] \\
& \text{handle}'_{\text{Nondet}} = \text{fold } (\text{alg}'_{\text{Nondet}} \nabla \text{con}) \ \text{gen}'_{\text{Nondet}} \\
& \text{gen}'_{\text{Nondet}} :: \text{TermMonad } m \ g \Rightarrow a \rightarrow m \ [a] \\
& \text{gen}'_{\text{Nondet}} \ x = \text{var } [x] \\
& \text{alg}'_{\text{Nondet}} :: \text{TermMonad } m \ g \Rightarrow \text{Nondet } (m \ [a]) \rightarrow m \ [a] \\
& \text{alg}'_{\text{Nondet}} \ (\text{Or } ml_1 \ ml_2) = \\
& \quad \text{do } \{ l_1 \leftarrow ml_1; \ l_2 \leftarrow ml_2; \ \text{var } (l_1 \ ++ \ l_2) \}
\end{aligned}$$

Notice that not much change has been necessary. We have generalized away from *Free g* into a type *m* that is constrained by *TermMonad m g*.

3.2 Parametricity: Fold/Build Fusion for Free

Hinze et al. [6] state that we get fold/build fusion for free from the *free theorem* [22,21] of the builder's polymorphic type. Hence, let us consider what the new type of *handler₁* buys us.

Theorem 1. *Assume that F_1, F_2, H_1, G_1 and A are closed types, with F_1, F_2 and H_1 functors. Given a function h of type $\forall m. (\text{TermMonad } m \ F_2) \Rightarrow \text{Free } F_1 \ A \rightarrow H_1 \ (m \ (G_1 \ A))$, two term monads M_1 and M_2 and a term monad morphism $\alpha :: \forall a. M_1 \ a \rightarrow M_2 \ a$, then:*

$$\text{fmap } \alpha \cdot h_{M_1} = h_{M_2} \quad (6)$$

where the subscripts of h denote the instantiations of the polymorphic type variable m .

If *handler₂* is a term monad morphism, then we can use the free theorem to determine *pipeline₁₂* in one step, starting from Equation (3).

$$\begin{aligned}
& \text{fmap } \text{handler}_2 \cdot \text{handler}_1 = \text{pipeline}_{12} \\
& \equiv \{ \text{Parametricity (6), assuming } \text{handler}_2 \text{ is a term monad morphism } \} \\
& \text{handler}_1 = \text{pipeline}_{12}
\end{aligned}$$

Unfortunately, *handler₂* is not a term monad morphism for the simple reason that H_2 is just an arbitrary functor that does not necessarily have a monad structure. Hence, in general H_2 is only term algebra.

instance *TermAlgebra* $H_2 \ F_2$ **where**
 $\text{var} = \text{gen}_2$
 $\text{con} = \text{alg}_2$

Fortunately, we can turn any term algebra into a term monad, thanks to the codensity monad, which is what we explore in the next section.

3.3 Codensity: *TermMonads* from *TermAlgebras*

The Codensity Monad It is well-known that the codensity monad *Cod* turns any (endo)functor *h* into a monad (in fact, *h* need not even be a functor at all). It simply instantiates the generalised monoid of endomorphisms (e.g., see [16]) in the category of endofunctors.

```
newtype Cod h a = Cod { unCod :: ∀x . (a → h x) → h x }
instance Monad (Cod h) where
  return x = Cod (λk → k x)
  Cod m >>= f = Cod (λk → m (λa → unCod (f a) k))
```

TermMonad Construction Given any term algebra *h* for functor *f*, we have that *Cod h* is also a term algebra for *f*.

```
instance TermAlgebra h f ⇒ TermAlgebra (Cod h) f where
  var = return
  con = algCod con
  algCod :: Functor f ⇒ (∀x . f (h x) → h x) → (f (Cod h a) → Cod h a)
  algCod alg op = Cod (λk → alg (fmap (λm → unCod m k) op))
```

Moreover, *Cod h* is also a term monad for *f*, even if *h* is not.

```
instance TermAlgebra h f ⇒ TermMonad (Cod h) f
```

The definition of *var* makes it easy to see that it satisfies the first term monad law in Equation (4). The proof for the second law, Equation (5) is less obvious:

$$\begin{aligned}
& \text{con op} \gg= f \\
\equiv & \{ \text{unfold } (\gg=) \} \\
& \text{Cod } (\lambda k \rightarrow \text{unCod } (\text{con op}) (\lambda a \rightarrow \text{unCod } (f a) k)) \\
\equiv & \{ \text{unfold con} \} \\
& \text{Cod } (\lambda k \rightarrow \text{unCod } (\text{alg}_{\text{Cod}} \text{con op}) (\lambda a \rightarrow \text{unCod } (f a) k)) \\
\equiv & \{ \text{unfold alg}_{\text{Cod}} \} \\
& \text{Cod } (\lambda k \rightarrow \text{unCod } (\text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow \text{unCod } m k) \text{ op}))) (\lambda a \rightarrow \text{unCod } (f a) k)) \\
\equiv & \{ \text{apply } \text{unCod} \cdot \text{Cod} = \text{id} \} \\
& \text{Cod } (\lambda k \rightarrow (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow \text{unCod } m k) \text{ op})) (\lambda a \rightarrow \text{unCod } (f a) k)) \\
\equiv & \{ \text{apply } \beta\text{-reduction} \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow \text{unCod } m (\lambda a \rightarrow \text{unCod } (f a) k)) \text{ op})) \\
\equiv & \{ \text{apply } \beta\text{-expansion} \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow (\lambda k \rightarrow \text{unCod } m (\lambda a \rightarrow \text{unCod } (f a) k)) k) \text{ op})) \\
\equiv & \{ \text{apply } \text{unCod} \cdot \text{Cod} = \text{id} \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow \text{unCod } (\text{Cod } (\lambda k \rightarrow \text{unCod } m (\lambda a \rightarrow \text{unCod } (f a) k)))) k) \text{ op})) \\
\equiv & \{ \text{fold } (\gg=) \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow \text{unCod } (m \gg= f) k) \text{ op})) \\
\equiv & \{ \text{apply } \beta\text{-expansion} \}
\end{aligned}$$

$$\begin{aligned}
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow (\lambda m \rightarrow \text{unCod } m \ k) (m \gg k)) \text{ op})) \\
\equiv & \quad \{ \text{apply } \beta\text{-expansion} \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow (\lambda m \rightarrow \text{unCod } m \ k) ((\lambda m \rightarrow m \gg k) \ m)) \text{ op})) \\
\equiv & \quad \{ \text{fold } (\cdot) \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow ((\lambda m \rightarrow \text{unCod } m \ k) \cdot (\lambda m \rightarrow m \gg k)) \ m) \text{ op})) \\
\equiv & \quad \{ \text{apply } \eta\text{-reduction} \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } ((\lambda m \rightarrow \text{unCod } m \ k) \cdot (\lambda m \rightarrow m \gg k)) \text{ op})) \\
\equiv & \quad \{ \text{apply } \text{fmap}\text{-fission and unfold } (\cdot) \} \\
& \text{Cod } (\lambda k \rightarrow \text{con } (\text{fmap } (\lambda m \rightarrow \text{unCod } m \ k) (\text{fmap } (\lambda m \rightarrow m \gg k) \text{ op}))) \\
\equiv & \quad \{ \text{fold } \text{alg}_{\text{Cod}} \} \\
& \text{alg}_{\text{Cod}} \text{ con } (\text{fmap } (\lambda m \rightarrow m \gg k) \text{ op}) \\
\equiv & \quad \{ \text{fold } \text{con} \text{ and } \eta\text{-reduce} \} \\
& \text{con } (\text{fmap } (\gg f) \text{ op})
\end{aligned}$$

3.4 Shifting to Codensity

Now we can write handler_2 as the composition of a term monad morphism $\text{handler}'_2$ with a post-processing function $\text{runCod } \text{gen}_2$:

$$\begin{aligned}
& \text{handler}_2 :: \text{Free } F_2 \ a \rightarrow H_2 \ a \\
& \text{handler}_2 = \text{runCod } \text{gen}_2 \cdot \text{handler}'_2 \\
& \text{handler}'_2 :: \text{Free } F_2 \ a \rightarrow \text{Cod } H_2 \ a \\
& \text{handler}'_2 = \text{fold } (\text{alg}_{\text{Cod}} \ \text{alg}_2) \ \text{var} \\
& \text{runCod} :: (a \rightarrow f \ x) \rightarrow \text{Cod } f \ a \rightarrow f \ x \\
& \text{runCod } g \ m = \text{unCod } m \ g
\end{aligned}$$

This decomposition of handler_2 hinges on the following property:

$$\text{fold } \text{alg}_2 \ \text{gen}_2 = \text{runCod } \text{gen}_2 \cdot \text{fold } (\text{alg}_{\text{Cod}} \ \text{alg}_2) \ \text{var} \quad (7)$$

This equation follows from the second fusion law for folds (2), provided that:

$$\begin{aligned}
& \text{gen}_2 = \text{runCod } \text{gen}_2 \cdot \text{var} \\
& \text{runCod } \text{gen}_2 \cdot \text{alg}_{\text{Cod}} \ \text{alg}_2 = \text{alg}_2 \cdot \text{fmap } (\text{runCod } \text{gen}_2)
\end{aligned}$$

The former holds as follows:

$$\begin{aligned}
& \text{runCod } \text{gen}_2 \cdot \text{var} \\
\equiv & \quad \{ \text{unfold } \cdot \} \\
& \lambda x \rightarrow \text{runCod } \text{gen}_2 \ (\text{var } x) \\
\equiv & \quad \{ \text{unfold } \text{runCod} \text{ and } \text{var} \} \\
& \lambda x \rightarrow \text{unCod } (\text{Cod } (\lambda k \rightarrow k \ x)) \ \text{gen}_2 \\
\equiv & \quad \{ \text{apply } \text{runCod} \cdot \text{Cod} = \text{id} \} \\
& \lambda x \rightarrow (\lambda k \rightarrow k \ x) \ \text{gen}_2 \\
\equiv & \quad \{ \beta\text{-reduction} \} \\
& \lambda x \rightarrow \text{gen}_2 \ x
\end{aligned}$$

$$\equiv \{ \eta\text{-reduction} \} \\ \text{gen}_2$$

and the latter:

$$\begin{aligned} & \text{runCod gen}_2 \cdot \text{alg}_{\text{Cod}} \text{alg}_2 \\ \equiv & \{ \text{unfold } \cdot \} \\ & \lambda op \rightarrow \text{runCod gen}_2 (\text{alg}_{\text{Cod}} \text{alg}_2 op) \\ \equiv & \{ \text{unfold runCod and alg}_{\text{Cod}} \} \\ & \lambda op \rightarrow \text{unCod} (\text{Cod} (\lambda k \rightarrow \text{alg}_2 (\text{fmap} (\lambda m \rightarrow \text{unCod } m \ k) op))) \text{gen}_2 \\ \equiv & \{ \text{apply runCod} \cdot \text{Cod} = \text{id} \} \\ & \lambda op \rightarrow (\lambda k \rightarrow \text{alg}_2 (\text{fmap} (\lambda m \rightarrow \text{unCod } m \ k) op)) \text{gen}_2 \\ \equiv & \{ \beta\text{-reduction} \} \\ & \lambda op \rightarrow \text{alg}_2 (\text{fmap} (\lambda m \rightarrow \text{unCod } m \ \text{gen}_2) op) \\ \equiv & \{ \text{fold runCod} \} \\ & \lambda op \rightarrow \text{alg}_2 (\text{fmap} (\lambda m \rightarrow \text{runCod gen}_2 \ m) op) \\ \equiv & \{ \eta\text{-reduction} \} \\ & \lambda op \rightarrow \text{alg}_2 (\text{fmap} (\text{runCod gen}_2) op) \\ \equiv & \{ \text{fold } \cdot \} \\ & \text{alg}_2 \cdot \text{fmap} (\text{runCod gen}_2) \end{aligned}$$

3.5 Fusion at Last

Finally, instead of fusing $\text{fmap handler}_2 \cdot \text{handler}_1$ we can fuse $\text{fmap handler}'_2 \cdot \text{handler}_1$ using the free theorem. This yields:

$$\text{pipeline}'_{12} = \text{fold alg}_1 \text{gen}_1$$

Now we can calculate the original fusion:

$$\begin{aligned} & \text{fmap handler}_2 \cdot \text{handler}_1 \\ \equiv & \{ \text{decomposition of handler}_2 \} \\ & \text{fmap} (\text{runCod gen}_2 \cdot \text{handler}'_2) \cdot \text{handler}_1 \\ \equiv & \{ \text{fmap fission} \} \\ & \text{fmap} (\text{runCod gen}_2) \cdot \text{fmap handler}'_2 \cdot \text{handler}_1 \\ \equiv & \{ \text{free theorem} \} \\ & \text{fmap} (\text{runCod gen}_2) \cdot \text{handler}_1 \end{aligned}$$

In other words, the fused version can be defined as:

$$\text{pipeline}_{12} = \text{fmap} (\text{runCod gen}_2) \cdot \text{fold alg}_1 \text{gen}_1$$

Observe that this version only performs a single *fold* and does not allocate any intermediate tree.

3.6 Repeated Fusion

Often a sequence of handlers is not restricted to two. Fortunately, we can easily generalize the above to a pipeline of n handlers

$$fmap^n handler_n \cdot \dots \cdot fmap handler_1 \cdot handler_0$$

where $handler_i = fold\ alg_i\ gen_i$ ($i \in 1..n$). This pipeline fusion can start by arbitrarily fusing two consecutive handlers $handler_j$ and $handler_{j+1}$ using the above approach, and then incrementally extending the fused kernel on the left and the right with additional handlers. These two kinds of extensions are explained below.

Fusion on the Right Suppose that $fmap handler_2 \cdot handler_1$ is composed with another handler on the right:

$$\begin{aligned} handler_0 &:: (TermMonad\ m_1\ F_1) \Rightarrow Free\ F_0\ a \rightarrow H_0\ (m_1\ (G_0\ a)) \\ handler_0 &= fold\ alg_0\ gen_0 \end{aligned}$$

to form the pipeline:

$$pipeline_{012} = fmap (fmap handler_2 \cdot handler_1) \cdot handler_0$$

Can we perform the fusion twice to obtain a single *fold* and eliminate both intermediate trees? Yes! The first fusion, as before yields:

$$pipeline_{012} = fmap (fmap (runCod gen_2) \cdot handler_1) \cdot handler_0$$

Applying *fmap* fission and regrouping, we obtain:

$$pipeline_{012} = fmap (fmap (runCod gen_2)) \cdot (fmap handler_1 \cdot handler_0)$$

Now the right component is another instance of the binary fusion problem, which yields:

$$pipeline_{012} = fmap (fmap (runCod gen_2)) \cdot fmap (runCod gen_1) \cdot handler_0$$

Fusion on the Left Suppose that $handler_2$ has the more specialised type:

$$handler_2 :: (TermMonad\ m_3\ F_3) \Rightarrow Free\ F_2\ a \rightarrow H_2\ (m_2\ (G_2\ a))$$

then we can compose $fmap handler_2 \cdot handler_1$ on the left with another handler:

$$\begin{aligned} handler_3 &:: Free\ F_3\ a \rightarrow H_3\ a \\ handler_3 &= fold\ alg_3\ gen_3 \end{aligned}$$

This yields a slightly more complicated fusion scenario:

$$pipeline_{123} = fmap (fmap handler_3 \cdot handler_2) \cdot handler_1$$

Of course, we can first fuse *handler*₂ and *handler*₃. That would yield an instance of fusion on the right. However, suppose we first fuse *handler*₁ and *handler*₂, after applying *fmap* fission.

$$pipeline_{123} = fmap (fmap handler_3) \cdot fmap (runCod gen_2) \cdot handler_1$$

Now we can shift the carrier of *handler*₃ to codensity and invoke the free theorem on *fmap* (*runCod* *gen*₂) · *handler*₁. This accomplishes the second fusion.

$$pipeline_{123} = fmap (fmap (runCod gen_3)) \cdot fmap (runCod gen_2) \cdot handler_1$$

Summary An arbitrary pipeline of the form:

$$fmap^n handler_n \cdot \dots \cdot fmap handler_1 \cdot handler_0$$

where *handler*_{*i*} = *fold alg_i gen_i* (*i* ∈ 1..*n*) fuses into

$$fmap^n (runCod g_n) \cdot \dots \cdot fmap (runCod gen_1) \cdot handler_0$$

3.7 Fusion all the Way

We are not restricted to fusing handlers, but can fuse all the way, up to and including the expression that builds the initial AST and to which the handlers are applied. Consider for example the *coin* example of Section 2.1. The free theorem of *coin*'s type is a variant of Theorem 1:

$$\alpha \text{ coin} = \text{coin}$$

where $\alpha :: \forall a . M_1 \ a \rightarrow M_2 \ a$ is a term monad morphism between any two term monads *M*₁ and *M*₂. We can use this to fuse *handle_{Nondet} coin* into *runCod gen_{Nondet} coin*. Of course this fusion interacts nicely with the fusion of a pipeline of handlers.

4 Pragmatic Implementation and Evaluation

This section turns the fusion approach into a practical Haskell implementation and evaluates the performance improvement.

4.1 Pragmatic Implementation

Before we can put the fusion approach into practice, we need to consider a few pragmatic implementation aspects.

Inlining with Typeclasses In a lazy language like Haskell, fusion only leads to a significant performance gain if it is performed statically by the compiler and combined with inlining. In the context of the GHC compiler, the inlining requirement leaves little implementation freedom: GHC is rather reluctant to inline in general recursive code. There is only one exception: GHC is keen to create type-specialised copies of (constraint) polymorphic recursive definitions and to inline the definitions of typeclass methods in the process.

In short, if we wish to get good speed-ups from effect handler fusion, we need to make sure that the effectful programs are polymorphic in the term monad and that all the algebras are held in typeclass instances. For this reason, all handlers should be made instances of *TermAlgebra*.

Explicit Carrier Functors The carrier functor of the compositional state handler is $s \rightarrow m_2-$. From the category theory point of view, this is clearly a functor. However, it is neither an instance of the Haskell *Functor* typeclass nor can it be made one in this syntactically invalid form. A new type needs to be created that instantiates the functor typeclass:

```
newtype StateCarrier s m a = SC { unSC :: s → m a }
instance Functor m ⇒ Functor (StateCarrier s m) where
  fmap f x = SC (fmap (fmap f) (unSC x))
instance TermMonad m2 f
  ⇒ TermAlgebra (StateCarrier s m2) (State s + f) where
  con = SC · (alg' State ∇ con State) · fmap unSC
  var = SC · gen' State
  gen' State :: TermMonad m f ⇒ a → (s → m a)
  gen' State x = const (var x)
  alg' State :: TermMonad m f ⇒ State s (s → m a) → (s → m a)
  alg' State (Put s' k) s = k s'
  alg' State (Get k) s = k s s
```

Now the following function is convenient to run a fused state handler.

```
runStateC :: TermMonad m2 f ⇒ Cod (StateCarrier s m2) a → (s → m2 a)
runStateC = unSC · runCod var
```

Unique Carrier Functors Even though the logging state handler has ostensibly⁴ the same carrier $s \rightarrow m_2-$ as the regular compositional state handler, we cannot reuse the same functor. The reason is that in the typeclass-based approach the carrier functor must uniquely determine the algebra; two different typeclass instances for the same type are forbidden. Hence, we need to write a new set of definitions as follows:

```
newtype LogStateCarrier s m a = LSC { unLSC :: s → m a }
```

⁴ The typeclass constraints on m_2 are different.


```

instance Functor m  $\Rightarrow$  Functor (LogStateCarrier s m) where
  fmap f x = LSC (fmap (fmap f) (unLSC x))
instance TermMonad m2 (Writer String + Void)
 $\Rightarrow$  TermAlgebra (LogStateCarrier s m2) (State s) where
  var = LSC · gen' State
  con = LSC · algLogState · fmap unLSC
runLogStateC :: TermMonad m2 (Writer String + Void)
 $\Rightarrow$  Cod (LogStateCarrier s m2) a  $\rightarrow$  (s  $\rightarrow$  m2 a)
runLogStateC = unLSC · runCod var

```

4.2 Evaluation

To evaluate the impact of fusion we consider several benchmarks implemented in different ways: running handlers over the inductive definition of the free monad (FREE) and to its Church encoding (CHURCH), the fully fused definitions (FUSED), and the conventional definitions of the state monad from MTL (MTL).

The benchmarks are run in the Criterion benchmarking harness using GHC 7.10.1 on a MacBook Pro with a 3 GHz Intel Core i7 processor, 16 GB RAM and Mac OS 10.10.3. All values are in milliseconds, and show the ordinary least-squares regression of the recorded samples; the R^2 goodness of fit is above 0.99 in all instances.

Benchmark	FREE	CHURCH	FUSED	MTL
<i>count</i> ₁				
10 ⁷	1,017	1,311	3	3
10 ⁸	10,250	13,220	29	29
10 ⁹	103,000	129,700	291	295
<i>count</i> ₂				
10 ⁶	684	746	167	213
10 ⁷	6,937	7,344	1,740	2,157
10 ⁸	102,700	98,010	17,220	20,300
<i>count</i> ₃				
10 ⁶	559	555	166	205
10 ⁷	5,759	5,561	1,618	2,132
10 ⁸	110,900	94,760	16,300	20,120
<i>grammar</i>	794	763	6	77
<i>pipes</i>	1,325	1,351	43	N/A

The *count*₁ benchmark consists of the simple count-down loop used by Kammar et al. [9].

```

count1 =
  do i ← get
    if i ≡ 0 then return i
    else put (i - 1); count1

```

We have evaluated this program with three different initial states of 10^7 , 10^8 and 10^9 . The results show that all representations scale linearly with the size of the benchmark. However, the fused representation is about 300 times faster than the free monad representations and matches the performance of the traditional monad transformers.

The *count₂* benchmarks extends the *count₁* benchmark with a *tell* operation from the *Writer* effect in every iteration of the loop. It is run by sequencing the state and writer handler. The improvement due to fusion is now much less extreme, but still quite significant.

The *count₃* benchmark is the *count₁* program, but run with the logging state handler that delegates to the writer handler. The runtimes are slightly better than those of *count₂*.

The *grammar* benchmark implements a simpler parser by layering the state and non-determinism effects. Again fusion has a tremendous impact, even considerably outperforming the MTL implementation.

The *pipes* benchmark consists of the simple producer-consumer pipe used by Kammar et al. [9]. We can see that fusion provides a significant improvement over either free monad representation. There is no sensible MTL implementation to compare with for this benchmark.

The results (in ms) show that the naive approaches based on intermediate trees, either defined inductively or by Church encoding, incur a considerable overhead compared to traditional monads and monad transformers. Yet, thanks to fusion they can easily compete or even slightly outperform the latter.

5 Related Work

5.1 Fusion

Fusion has received much attention, and was first considered as the elimination of trees by Wadler [23] with his so-called deforestation algorithm, which was then later generalized by Chin [3].

From the implementation perspective, Gill et al. first introduced the notion of shortcut fusion to Haskell [4], thus allowing programs written as folds over lists to be fused. Takano and Meijer showed how fusion could be generalized to arbitrary datastructures [19]. The technique of using free theorems to explain certain kinds of fusion was applied by Voigtländer [20].

Work by Hinze et al. [6] builds on recursive coalgebras to show the theory and practice of fusion, limited to the case of folds and unfolds. Later work by Harper [5] provides a pragmatic discussion, that bridges the gap between theory and practice further, by discussing the implementation of fusion in GHC with inline rules. Harper also considers the fusion of Church encodings.

More recently, recursive coalgebras appear in work by Hinze et al. [7] where conjugate hylomorphisms are introduced as a means of unifying all presently known structured recursion schemes. There, the theory behind fusion for general datatypes across all known schemes is described as arising from naturality laws of adjunctions. Special attention is drawn to fusion of the cofree comonad, which is the dual case of free monads we consider here.

5.2 Effect Handlers

Plotkin and Power were the first to explore effect operations [14], and gave an algebraic account of effects [15] and their combination [8]. Subsequently, Plotkin and Pretnar [13] have added the concept of handlers to deal with exceptions. This has led to many implementation approaches.

Lazy Languages Many implementations of effect handlers focus on the lazy language Haskell.

For the sake of simplicity and without regard for efficiency, Wu et al. [24] use the inductive datatype representation of the free monad for their exposition. They use the *Data Types à la Carte* approach [18] to conveniently inject functors into a co-product; that approach is entirely compatible with this paper. Wu et al. also generalize the free monad to *higher-order functors*; we expect that our fusion approach generalizes to that setting.

Kiselyov et al. [10] provide a Haskell implementation in terms of the free monad too. However, they combine this representation with two optimizations: 1) the codensity monad improves the performance of (\gg), and 2) their *Dynamic*-based open unions have a better time complexity than nested co-products. Due to the use of the codensity monad, this paper also benefits from the former improvement. Moreover, we believe that the latter improvement is unnecessary due to the specialisation and inlining opportunities that are exposed by fusion.

Van der Ploeg and Kiselyov [12] present an implementation of the free monad with good asymptotic complexity for both pattern matching and binding; unfortunately, the constant factors involved are rather high. This representation is mainly useful for effect handlers that cannot be easily expressed as folds, and thus fall outside of the scope of the current paper.

Behind a Template Haskell frontend Kammar et al. [9] consider a range of different Haskell implementations and perform a performance comparison. Their basic representation is the inductive datatype definition, with a minor twist: the functor is split into syntax for the operation itself and a separate continuation. This representation is improved with the codensity monad. Finally, they provide—without explanation—a representation that is very close to the one presented here; their use of the continuation monad instead of the codensity monad is a minor difference.

Both Atkey et al. [1] and Schrijvers et al. [17] study the interleaving of a free monad with an arbitrary monad, i.e., the combination of algebraic effect handlers and conventional monadic effects. We believe that our fusion technique can be adapted for optimizing the free monad aspect of their settings.

Strict Languages In the absence of lazy evaluation, the inductive datatype definition of the free monad is not practical.

Kammar et al. [9] briefly sketch an implementation based on delimited continuations. Schrijvers et al. [17] show the equivalence between a delimited continuations approach and the inductive datatype; hence the fusion technique presented in this paper is in principle possible. However, in practice, the codensity monad used for fusion is likely not efficient in strict languages. Hence, effective fusion for strict languages remains to be investigated.

5.3 Monad Transformers

Monad transformers, as first introduced by Liang et al. [11], pre-date algebraic effect handlers as a means for modelling compositional effects. Yet, there exists a close connection between both approaches: monad transformers are fused forms of effect handlers. What is particular about their underlying effect handlers is that their carriers are term algebras with a monadic structure, i.e., term monads. This means that the *Cod* construction is not necessary for fusion.

6 Conclusion

We have explained how to fuse algebraic effect handlers by shifting perspective from free monads to term monads. Our benchmarks show that, with a careful code setup in Haskell, this leads to good speed-ups compared to the free monad, and allows algebraic effect handlers to compete with the traditional monad transformers.

Acknowledgments

The authors would like to thank Ralf Hinze for suggesting that they consider the specification that free monads satisfy in terms of adjunctions, and to James McKinna who helped baptize *con* for constructing terms. We are also grateful for the feedback from the members of the IFIP Working Group 2.1 on an early iteration of this work. This work has been partially funded by the Flemish Fund for Scientific Research (FWO).

References

1. Atkey, R., Johann, P., Ghani, N., Jacobs, B.: Interleaving data and effects (2012), submitted for publication
2. Burris, S., Sankappanavar, H.P.: A Course in Universal Algebra. No. 78 in Graduate Texts in Mathematics, Springer-Verlag (1981)
3. Chin, W.N.: Safe fusion of functional expressions ii: Further improvements. *Journal of Functional Programming* 4, 515–555 (Oct 1994)
4. Gill, A., Launchbury, J., Peyton Jones, S.L.: A short cut to deforestation. In: *Proceedings of the Conference on Functional Programming Languages and Computer Architecture*. pp. 223–232. FPCA '93, ACM, New York, NY, USA (1993)

5. Harper, T.: A library writer's guide to shortcut fusion. In: Proceedings of the 4th ACM Symposium on Haskell. pp. 47–58. Haskell '11, ACM, New York, NY, USA (2011)
6. Hinze, R., Harper, T., James, D.: Theory and practice of fusion. In: Hage, J., Morazán, M. (eds.) Implementation and Application of Functional Languages, Lecture Notes in Computer Science, vol. 6647, pp. 19–37. Springer Berlin Heidelberg (2011)
7. Hinze, R., Wu, N., Gibbons, J.: Conjugate hylomorphisms – the mother of all structured recursion schemes. In: Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 527–538. POPL '15, ACM, New York, NY, USA (2015)
8. Hyland, M., Plotkin, G.D., Power, J.: Combining effects: Sum and tensor. Theoretical Computer Science 357(1-3), 70–99 (2006)
9. Kammar, O., Lindley, S., Oury, N.: Handlers in action. In: Proceedings of the 18th ACM SIGPLAN International Conference on Functional programming. pp. 145–158. ICFP '14, ACM (2013)
10. Kiselyov, O., Sabry, A., Swords, C.: Extensible effects: an alternative to monad transformers. In: Proceedings of the 2013 ACM SIGPLAN symposium on Haskell. pp. 59–70. Haskell '13, ACM (2013)
11. Liang, S., Hudak, P., Jones, M.: Monad transformers and modular interpreters. In: Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 333–343. POPL '95, ACM, New York, NY, USA (1995)
12. Ploeg, A.v.d., Kiselyov, O.: Reflection without remorse: Revealing a hidden sequence to speed up monadic reflection. In: Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell. pp. 133–144. Haskell '14, ACM, New York, NY, USA (2014)
13. Plotkin, G.D., Matija, P.: Handling algebraic effects. Logical Methods in Computer Science 9(4) (2013)
14. Plotkin, G.D., Power, J.: Notions of computation determine monads. In: Nielsen, M., Engberg, U. (eds.) Foundations of Software Science and Computation Structures. Lecture Notes in Computer Science, vol. 2303, pp. 342–356. Springer (2002)
15. Plotkin, G.D., Power, J.: Algebraic operations and generic effects. Applied Categorical Structures 11(1), 69–94 (2003)
16. Rivas, E., Jaskelioff, M.: Notions of computation as monoids. CoRR abs/1406.4823 (2014), <http://arxiv.org/abs/1406.4823>
17. Schrijvers, T., Wu, N., Desouter, B., Demoen, B.: Heuristics entwined with handlers combined. In: Proceedings of the 16th International Symposium on Principles and Practice of Declarative Programming. pp. 259–270. PPDP '14, ACM (Sep 2014)
18. Swierstra, W.: Data types à la carte. Journal of Functional Programming 18(4), 423–436 (2008)
19. Takano, A., Meijer, E.: Shortcut deforestation in calculational form. In: Proceedings of the Seventh International Conference on Functional Programming Languages and Computer Architecture. pp. 306–313. FPCA '95, ACM, New York, NY, USA (1995)
20. Voigtländer, J.: Proving correctness via free theorems: The case of the destroy/build-rule. In: Proceedings of the 2008 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation. pp. 13–20. PEPM '08, ACM, New York, NY, USA (2008)
21. Voigtländer, J.: Free theorems involving type constructor classes: Functional pearl. In: Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming. pp. 173–184. ICFP '09, ACM, New York, NY, USA (2009)

- 22. Wadler, P.: Theorems for free! In: Proceedings of the Fourth International Conference on Functional Programming Languages and Computer Architecture. pp. 347–359. FPCA '89, ACM, New York, NY, USA (1989)
- 23. Wadler, P.: Deforestation: Transforming programs to eliminate trees. Theoretical Computer Science 73(2), 231–248 (Jan 1990)
- 24. Wu, N., Schrijvers, T., Hinze, R.: Effect handlers in scope. In: Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell. pp. 1–12. Haskell '14, ACM, New York, NY, USA (2014)